

This article offers an unfiltered look at a large cross section of the different rapid prototyping technologies available today; from a guy with one of the biggest RP toy boxes in the world as the manager of the Rapid Prototyping Laboratory at NASA's Marshall Space Flight Center (MSFC) in Huntsville, AL, USA. NASA's current operation capacity is nine RP machines, representing eight actual technologies. The article presents a realistic, unbiased look at the technologies and offers advice on what to do and where to go for the best solution to your rapid prototyping needs.

RP That's Right For You!

Ken Cooper

So, you've seen the advertisements, the success stories, and the glam of this "rapid prototyping" (RP) technology. Building parts from the inside out, fabricating "unmachineable" geometry, saving 50-70% on time and money\$. . .but does it really work?

In a nutshell, "Yes."

Which leads on to: Which system is the best?

The answer is again, "Yes".

First off, you may be thinking, "Why so many processes under one roof?" Although the Marshall Space Flight Center (MSFC) particularly looks at the technologies for applications in microgravity and outer space, one of the primary goals of the lab that has evolved over the last decade is to evaluate each technology and identify ideal applications for each within NASA not only for research purposes, but also for standard fabrication best practices. The goal is to maintain a position on the cutting edge (often the "bleeding edge") of the technology, and then disseminate the information out to industry. It gives interested individuals a fresh, unbiased look at how RP technologies can benefit their particular operation. In doing so, it was found that there are no "best" RP technologies for any given application, but more often several that apply depending on the customer's constraints. That is most likely the reason that there are still several different technologies surviving the marketplace, which is great for customers who prefer to have choices when shopping for what they need.

Keep in mind that RP is still an "infant" technology, compared with standard machining, casting, etc. Also, innovative end users are ultimately the ones who determine what "niches" the technologies fall in to, and they will obviously vary from site to site. The eight processes evaluated here are as follows:

1. Stereolithography (SL)
2. Multi Jet Modelling (MJM)
3. 3D Printing
4. Selective Laser Sintering (SLS)
5. Laminated Object Manufacturing (LOM)
6. Fused Deposition Modeling (FDM)
7. Droplet Based Manufacturing
8. Laser Engineered Net Shaping (LENS).

Data and rankings come from internal testing as well as polling of users in the field. With all that said, let's cut to the chase, shall we?

RP 101

In case you are reading this and are not familiar with RP technology in general, here is the basic principle: A computer-aided design (CAD) model is used to generate virtual "slices" of the desired model, which are then downloaded to an RP machine to fabricate. All of the processes work on the same additive technique, where the layers are added from the bottom upwards to essentially "grow" the part. The actual steps required to get a model from an RP device once the CAD data has been received generally lie in four categories:

1. Front-end software setup.
2. Physical hardware setup.
3. Machine build operation (unattended).
4. Post processing.

The fourth step is the most labour intensive step by far, where the physical part is removed from the RP machine and traditional hand modelling skills are adopted to finish the product. As most of the up-front software processing is now becoming quite "user friendly", the last three categories will be used to evaluate each system in addition to applications, general maintenance, and operating costs. The technologies herein are not listed in any particular order or rank, and the ones that have fallen asunder are not covered for brevity's sake.

Stereolithography (SL)

SL is probably the name most synonymous with RP, as it was indeed the first RP technology to hit the streets. SL remains one of the more prominent systems in the industry, with a large market share and quality equipment. The technology itself, while not changing in concept, has improved significantly in the past decade. The concept within relies on actively curing photo-sensitive polymers with a low-powered laser (25-150 mW) to form the desired shape. Most of the resins used today are epoxy based, and all newer systems operate using solid-state lasers. Even still, a large part of the market is relying on the older gas-based lasers, which were originally used to kick off the technology around twelve years ago.

The physical hardware setup requires proper handling of the liquid based photopolymer materials, but otherwise consists simply of inserting a clean build grate and then pouring extra resin into the vat until the machine "beeps" at you to stop. Newer systems actually have an added reservoir to automatically keep the resin at the proper level for building.

Actual unattended machine build time is the longest link in the cycle, as parts may sometimes build for several days to a week if they are very large. The "speed" of the build depends primarily on the density of the parts, due to the raster scan curing technique of the system. For example, a hollow cylinder will build much faster

than a solid cylinder of the same size.

Post processing is a relatively laborious process, although many users have perfected the art. Parts must be drained and chemically cleaned of excess resin prior to support removal, hand finishing, and then final ultraviolet curing. Support structures, which are basically expendable ribs that are built along with the part to support overhangs in the part, must be physically removed by hand tools as they are built of the parent material. Fortunately, smart techniques have been developed that allow the supports to build very fine and web-like structures to aid in easier removal. Parts that are built for investment casting patterns must go through additional steps to draw excess resin from internal voids within the part shell. This includes draining, centrifuging, suction, and then re-sealing the resin exit points. Keeping the machines clean are a chore, but a must-do.

For most geometry, SL ranks about third or fourth in terms of build speed and probably first in accuracy and surface quality. Facility requirements are somewhat extensive for initial setup, however multiple machines can share the same post processing equipment. Many applications are viable from small, intricate, delicate features to functional tooling and test pieces. The larger frame SL systems are of the most expensive to purchase on the market. And although customers may get initial sticker shock from the annual maintenance contract costs, a poll of users turned up that it is actually well worth the money spent. [1]

Multi Jet Modelling

From the makers of SL comes Multi Jet Modeling (MJM), an "office modeller" system that prints molten polymer from arrays of ink jets to form a waxy-like part. Kudos to easy setup - with automatic part placement users do not even have to view the STL file, with just a few mouse clicks parts can start building. The machines are networked with IP addresses, so parts can be downloaded and queued from multiple remote locations.

Physical hardware set up consists of inserting a build plate into the machine until it locks, closing the door and pressing the "online" button. Occasional maintenance requires that you remove a tray of waste wax from jet purges, as well as replacing the build material bottle by taking the lid off and dropping it down the hatch.

The unattended machine build sequence is fairly quick, with full envelope parts not taking more than 24-28 hours, and most builds within a working day. The number or size of parts in the XY plane (length and width) really does not matter; only the Z axis, or height, of the parts essentially determines the build time. This provides for great economy of scale on multiples of small parts.

Post processing is somewhat delicate; as the parts are basically wax, so care must be taken in removing the support structures. Relatively speaking, it is still quite simple, and can be done with minimal tools. A good trick is to chill the part in a freezer for about half an hour to make the supports brittle (they are basically small wax needles), and then rub the majority of them off with gloved fingers

and a toothbrush. Unlike some others, the machine has no tolerance for any angle of overhang; therefore there are often a lot of support structures to clean out and hence a good bit of wasted material.

The MJM system ranks about third in build speed for most parts, and second only to stereolithography in surface finish although down-facing surfaces exposed to supports have to be hand finished. Applications primarily rest in concept modelling, as dimensional accuracies are lower and the materials are not very durable. Another application for which the waxy material lends itself is the investment-casting sector, as the materials melt only about 20°C higher than traditional foundry waxes. Facility requirements are minimal, needing only room for a photocopier-sized machine and house current with a network connection. Build material tends to be quite expensive, especially factoring in the high waste generation. The machines however currently have one of the lowest sticker prices on the market, as well as a very reasonable maintenance contract. [1]

Three Dimensional Printing

Three Dimensional Printing (3DP) was developed at MIT and is applied in three machine formats: concept modelling, direct metals, and ceramics. This review covers the concept modeling system. Also sold as an "office modeller", this technology uses multiple ink jets to deposit a liquid binder onto a powder bed, in essence printing the cross section into the powder layer by layer. The two materials currently available are starch-based and gypsum-based; the former being rougher in surface but more amenable to investment casting patterns and the latter being smoother and stronger and hence more recommended for concept verification modelling. Software setup is fairly easy, and comes with a nice CD-ROM training video to take you step-by-step through the process - something this author would like to see with every RP system.

The machine setup can tend to be a bit on the messy side, as users are handling powders and must take that in to consideration. Fortunately the "startup kit" includes a vacuum cleaner to keep it under control. Frequent ink-jet replacement has been a common complaint in the past, although the vendor has increased the life of its cartridges now (by about twofold) and made the jets much easier to change. Binder bottles usually last for several months and are easy to change out, and the overflow bottles last a very long time. The vacuum bags to capture powder overflow must be changed frequently but are commercially available items and easy to change.

The machine build time is the claim to fame of this process, as it still holds the record for the fastest unattended build time of most geometry. The build time is based primarily on the Z axis height of the part(s), with an average of about one vertical inch per hour. Almost all part sizes can be completed in a single working day of eight to nine hours, with most of them in just an hour or two. The machine has a nice feature that will allow it to start at a predetermined time during the night so it will actually finish near the beginning of the morning shift, for example.

Post processing consists of several steps and can be quite tedious with delicate parts. The as-built parts are soft (the starch more so

than the gypsum) and must be handled with care prior to infiltration. Fortunately, the powder bed eliminates the need for physical support structures, so no mechanical operations are required. Basically, the powder must be removed from around the part with the vacuum cleaner, and then it is placed in a small glove box to blow the remaining powder away with a nifty airbrush. Afterwards the parts must be infiltrated with the user's choice of material to give the part its final consistency. Infiltrants range from paraffin wax to flexible materials to very hard thermosets, resulting in a nice variety of applications. Most unused powder can be recycled; therefore material waste is not as big a problem as in some other processes.

The 3DP ranks first in build speed and overall part economy, and about seventh in surface finish and accuracy. The machine list price is very competitive with other office modelers, and the expendable materials costs and annual maintenance contracts are relatively low as well. The 3DP system now even boasts a full colour modelling capability, which should broaden the concept verification capability even more. [2]

Selective Laser Sintering (SLS)

The SLS process is a laser-based system that incorporates powders as build materials. SLS works by sweeping a thin layer of powder over the build area, and then melting the cross section in raster fashion with a CO2 laser (50-100W). About a dozen materials are available for use in SLS systems, including polymers, core sand, and even semi-metallic materials for tooling applications. Software setup requires a more skilled user, although once defaults are locked in on a material the learning curve shortens significantly.

The machine setup for a build requires handling powders and hence appropriate safety precautions, and can generate a good bit of dust and mess depending on the material being used. The build chamber must be kept under Nitrogen purge for most materials, and the vendor recommends resifting the powders and mixing 50/50 with virgin powder if you plan to recycle your unused material after a build.

The machine build time is typically very fast for most geometry, and is dependant mostly on the cross sectional density of the part and the Z axis height. If you import multiple parts and run them in the same batch, you can actually vary laser power per part in the build. This is a nice feature if for example a single part request gets cancelled during a batch build, you can simply turn the laser power to zero on that particular part and speed up the rest of the process. The slice-on-the-fly technology allows for other parameters (i.e. layer thickness, powder feed, temperature) to be changed at any given time during the build if necessary.

The post-processing phase again requires handling powders, as the part "cake" is removed from the chamber and taken to a break-out table to excavate the part. The table has a built-in air handler, which is a plus, to capture flying dust during part cleanup. Typical tools for cleanup are stiff brushes, dental picks, fingers, and compressed air. Sometimes, abrupt changes in cross section can cause curling, so a small amount of support material is usually added in the form of a large wafer at that cross section. These thin wafers

are trimmed off and discarded during cleanup. Excess powder from the cake runs through a sieve on the table surface and is captured for reuse. It's about a one to three hour time investment depending on the operator and the parts.

The SLS ranks about second in build speed, and is high on the dimensional tolerance scale as well. Accuracy and durability depend on the material being used, some of which receive high marks as well. Machines, maintenance, and expendables are all on the higher end of the cost scale, yet still competitive with comparative systems. Facility requirements are quite significant, as the machines are quite large and require dedicated powder and nitrogen gas hookups. Peripheral equipment purchased separately includes the break out table, powder sifter, and vacuum cleaner, not to mention a hydrogen furnace for the metal powders. Users will need a good bit of floor space and probably a dedicated operator for this system, but overall this system will achieve high-end results. [3]

Fused Deposition Modelling (FDM)

The FDM technique is an extrusion-based process, where a thin ribbon of molten polymer is extruded to lay down plastic "weld-beads" layer-by-layer in raster fashion to form the part. The most common analogy I've heard is a "hot-glue gun on a gantry." The FDM systems offer high temperature polymer materials and an easy, yet powerful front-end software interface.

Machine setup is quite simple, in that an expendable foam build plate is pinned into a tray that locks right into the machine. The sliced file is then ported to the machine from a PC (newer machines are networked), the head is positioned and purged, and then it can start building.

Build time is a bit slower than laser-based machines, mainly due to mechanical and inertial constraints of moving a large extrusion head at high speeds, not to mention the need for cohesion of the material. Regardless, significant speed improvements have been made within the technology over the last decade. If build material runs out during mid run, the machine is smart enough to pause itself and demand more before continuing, which has always been a strong feature. Part recovery after a power failure is also nice.

Post processing comes in two forms: mechanical and water-soluble. Mechanical uses a second material similar in composition to the build that is peeled away from the final part with picks and pliers. The difficulty of this operation depends entirely on the geometry of the part built. Internal structures and delicate features can ruin your day here. Fortunately, FDM has the capability (thanks to cohesion and surface tension of polymers) to build up to about a forty-five degree angle with support structures and can span sometimes over a quarter-inch or better as well. And luckily, the software is smart enough to make this determination up front and only put the supports where needed. When it comes to the water-soluble option, apparently a lot of the stress is removed in that the support material can be chemically removed in a bath and thus spare the small part features as well as your fingertips.

Typical part build time comes in around fifth, with high marks for durability and functionality and accuracy in the mid-range. Facility setup is a breeze on the smaller systems; basically the requirements are house current and a serial cable to your PC and you are running parts. Applications lean more to the functional prototyping side, although concept verification is still an economical option. The water-soluble option requires additional equipment for the cleanup operations. Machine costs range the spectrum, with smaller machines on the low end and larger, faster machines on the high end. Maintenance and material costs are fairly reasonable as well, and the machine fits in either an office or a shop floor environment. [4]

Laminated Object Manufacturing (LOM)

Once believed dead but now operating under a different company, LOM is a unique RP system in that it is more of a hybrid of additive and subtractive fabrication technology. Layers of adhesive paper are bonded with a heated roller, and then the part cross section is cut out with a laser. Internal voids in and between parts are then hashed up into cubes with the laser for later removal. The resulting parts have the qualities of an actual wood pattern.

The user software and machine setup are part of the same step. No pre-slicing is required as the LOM incorporates slice-on-the-fly technology, so the software is mainly used to put in the build parameters and for manual operations while setting up the machine for a build. Mirrors, lenses, and tracks exposed to the smoke and ash of the burned paper must be cleaned per part run. A metal build platen is placed in the machine, and then strips of double-backed foam tape are manually placed and cut out to match the part's maximum cross sectional bounding box. After several steps and about a half-hour or so, the actual part run can be started.

Part build time varies significantly with geometry and the user's tolerance for handwork after the build is done. The main dependence on speed is the XY (length and width) of the part, in addition to the Z height. Smaller, intricate parts typically require smaller crosshatched cubes, which in turn slow down each layer. The process is very fast, however, on large parts with dense cross sections. This is due to the fact that the material is deposited solid, and only the outlines are drawn with the laser as opposed to raster filling. Unfortunately paper jams sometimes occur, or the machine may run out of paper over night. It does have a feature to dial up your pager if an error occurs.

Post processing is perhaps the most labour intensive of any RP process discussed here. The part block is chiselled from the build platen, and then the cubes of void material must be removed with picks, pliers, knives, hammers—you get the point. Count on a few hours manual labour here, and keep plenty of Band-Aids around. Nevertheless, the resulting parts have an attractive wood-like appearance that might just make it worth the wait. The foam tape is removed from the platen with a mild citrus chemical solution and is then ready to build upon again. Parts are sealed in wood sealer and then can be treated pretty much the same as ordinary wood parts.

On small, intricate parts the LOM comes in about sixth in speed,

however for larger parts with dense cross sections it runs an easy first place. Processing of the parts ranks inversely proportional to the complexity of the pattern, and the material properties are basically wood-like. The surface quality can actually be pretty good, especially after sealing the material sands and finishes quite nicely. System cost is on the lower end, and the annual maintenance contract is in the mid-range. Facility setup requires standard power hookups, but a significant ventilation system (300 cfm duct fan) to pull smoke and fumes from the build envelope to the outdoors. [5]

Droplet Based Manufacturing (DBM)

Got Accuracy? DBM is yet another ink-jet based technology that deposits layers of waxy resin to form the part. A single jet deposits the build material in very fine layers, usually around 0.5mm, while a second jet is used to deposit a lower temperature solid support structure for overhangs. After each deposited layer a milling arm passes over the part to maintain Z height accuracy as well.

The machine setup requires purging and cleaning the jets, and placing an adhesive foam build pad onto a removable metal plate in the machine. Additionally, the machine uses rolls of "ticker tape" to check jet performance, which must be replaced. Loading the build material consists of pouring the wax pellets down the chute. The foam pad is milled a few times with the cutter to insure a flat build surface.

The building step is slower than any other due to the accuracy trade off of a very thin layer size, for example a one-half inch tall component is a 100-layer part on most RP machines, but is 1000 layers on this one. The system has a built-in jet check before and after each layer, and when a jet fails it will purge, mill off the previous layer, and build it again. The user must purge the jets daily, even if the part being built takes weeks or else the part starts doing the Charleston (one step forward and two steps back). The best prevention: build smaller parts. Otherwise, with a little TLC you can keep the machine running. A power failure recovery option is a default and necessary option here.

Post processing requires that the platen be heated on a hotplate and then the foam is peeled away. Sticky goo from the tape is scraped and chemically cleaned from the platen. The part and foam are then subjected to a warm chemical solution in a contained glove box for several hours to soften the support material and dissolve the foam. Afterwards the remaining support clinging to the part is soft and is gently removed by handwork. The cleaner solution is reused for a while, but then must be disposed of in an environmentally correct fashion.

Accuracy gets the top slot on this process, whereas speed comes in last place. Surface finish is good on both sides of the part, and supports are solid and removed chemically there, are no mechanical ridges underneath overhangs as is seen in other technologies. The majority of applications suited to this technology are very small parts, less than 2 or 3 inches each dimension that require tight dimensional stability. The machine does have a much larger working envelope if accuracy overrides the need for speed by a very large

margin. Material strength is consistent with other wax patterns, and the material lends itself to investment casting applications. System cost and maintenance are on the lower end, consistent with other office modellers. The smaller machine comes with a fairly small footprint, runs on house current and hooks directly to a PC, on which the relatively simple operating software is installed. [6]

Laser Engineered Net Shaping (LENS)

Need a little more than waxes and polymers? The LENS system applies additive fabrication to metals ranging from stainless steel to titanium. The metal is supplied in powder form by an argon gas stream into the focus point of a high-powered laser (500W-1.5kW), while tracing the shape of the part. A good analogy is "artistic welding", as tiny weld beads of the material are deposited side-by-side and layer-by-layer to form the part.

The machine setup requires bolting in a thick metal plate to build upon, usually of the same material to be deposited and around one-quarter inch thick. Metal powder is loaded into the feed hopper, and internal filters must be checked and replaced often. Argon gas is constantly supplied and recirculated through the isolated build chamber to maintain very low oxygen levels, on the order of a few parts per million. The laser is supplied with a constant flow of chilled water for coolant, and the coolant reservoir inside the laser power supplied must be refilled with deionized water before each build. The STL file is then converted to machine code and downloaded to the machine for building.

The part build sequence is not quite as automatic as other RP processes yet, as closed loop control is in the final phases of development. Some tweaking of powder flow rate and laser power is done manually throughout the build. The operation is semi-unattended, where once you dial in the parameters on a part it is usually safe to walk away for a few hours, but the machine usually is not left alone overnight just yet.

Post processing requires machine shop capabilities, as parts must be cut from a plate, which they are basically joined to completely, and then final machined to achieve a smooth finish and accuracy. Care must be taken to avoid thermal stresses in thicker parts, which can cause even the sturdiest build plate to buckle during the process.

LENS parts get the top slot for materials strength and durability, topping off even wrought metals of the same alloys. Parts have a somewhat surface finish and rank lower in accuracy, therefore most applications oversize the part and machine back down to spec. Speed, depending on the material, currently ranks about sixth or seventh compared to traditional RP processes. The LENS process is still new, even compared with the other RP technologies, and still under continuous development as well. With a handful of users (about 12 at last count), applications continue to emerge on a frequent basis, and I believe you'll be seeing a lot more of this technology in the future. [7]

Your Next Move

Again you ask, "So, which process is best?" Again I answer, "Yes." Re-phrase your question to "Which process is best for my application?" and you will get a straighter answer. My first suggestion is to reread this article and down select to a few technologies that sound like they meet your needs. Consider all of the criteria: system cost, facility requirements, ease of use, your typical timeline on projects, available man-power for post processing and system operation, strength requirements, tolerance and surface quality. Next, go to the RP&T Service Bureaux Review in this magazine or look on the Internet, pick some companies near you that use the technologies you chose, and have them build you some parts of each and put them to the test. You may find the system that's right for you to purchase, or you may never even buy your own machine if you build a good relationship with a local service bureau that has what you need.

The author would like to thank the MSFC rapid prototyping staff members Glenn Williams, Ashante Allen, OC Okike and Rob Minor for peer reviewing this article, and watering down some of my wittiness.

The author of the textbook Rapid Prototyping Technology: Selection and Application [8], Ken Cooper has been involved in the rapid prototyping industry for about nine years. He is the co-founder and current manager of the Rapid Prototyping Laboratory at NASA's Marshall Space Flight Center (MSFC) in Huntsville, AL, USA. In that time, he has acquired and operated at least ten different RP technologies (some of which were featured in a previous TCT article [9]). NASA's current operation capacity is nine RP machines, representing eight actual technologies. In addition, Ken also represents NASA on the National Center for Manufacturing Sciences' Rapid Prototyping Technology Advancement team (NCMS-RPTA), which is an ongoing industry/government collaboration effort to promote competitive development and increased application of RP technologies.

1. www.3dsystems.com
2. www.zcorp.com
3. www.dtm-corp.com
4. www.stratasys.com
5. www.cubictechnologies.com
6. www.solid-scape.com
7. www.optomec.com
8. Kenneth G. Cooper, "Rapid Prototyping Technology: Selection and Application" Marcel Dekker, Inc. January 2001, New York, 264 pages, illus. ISBN: 0-8247-0261-1
9. March 2001 TCT Article "Here LiesŠ" Obituaries of RP Manufacturers by Wohlers and Grimm